arXiv:1701.07058v2 [cs.GT] 3 Mar 2017

If you are not paying for it, you are the product:
How much do advertisers pay for your personal data?

Panagiotis Papadopoulos[1], Nicolas Kourtellis[2], Pablo Rodriguez Rodriguez[3], Nikolaos Laoutaris[2]

[1]FORTH-ICS, Greece, [2]Telefonica I+D, Spain, [3]Telefonica Alpha, Spain
[1]panpap@ics.forth.gr, [2,3]{name.lastname}@telefonica.com

Abstract

Online advertising is progressively moving towards a programmatic model in which ads are matched to actual interests of individuals collected as they navigate the web. Letting the huge debate around privacy aside, a very important question in this area, for which very little is known, is: *How much do advertisers pay for an individual's personal data?*

In this study, we develop a first of its kind methodology for computing exactly that – the price paid for one's privacy – and we do that in real time. Our approach is based on tapping on the Real Time Bidding (RTB) protocol to collect cleartext and encrypted prices for winning bids paid by advertisers in order to place targeted ads. Our main technical contribution is a method for tallying winning bids even when they are encrypted. We achieve this by training a model using as ground truth prices obtained by running our own "probe" ad-campaigns. We implement our methodology through a browser addon and a back-end server that provides it with fresh models for encrypted bids. We validate our methodology using a one year long trace of 1600 mobile users and demonstrate that it can estimate a user's worth with more than 82% accuracy.

1 Introduction

In today's data-driven economy, the amount of user data an IT company holds has a direct and non-trivial contribution to its overall market valuation [74]. Not surprisingly therefore, companies pursue with full strength the acquisition of ever more detailed personal data, either by collecting them while providing free services [3, 34], or by buying them from others [84, 10, 51]. Conservative estimates put the market of personal data to hundreds of billion dollars [23, 86, 18].

Digital advertising is one of the main means of monetising user data. It grew to $194.6 billion in 2016 [76] of which $108 billion were due to mobile advertising. In fact, more and more companies rush to participate in this rapidly growing advertising business either as advertisers, ad-networks, demand-side platforms, data management platforms, or all of the above. For these companies to increase their market share, they need to deliver more effective and highly targeted advertisements. A way to achieve this is through programmatic instantaneous auctions. An important enabler for this kind of auctions is the Real-Time Bidding (RTB) protocol for transacting digital display ads in real time. RTB has been growing with an annual rate of 128% [81], and currently accounts for 74% of programmatically purchased advertising. In US alone it created a revenue of $8.7 billion in 2016 [7].

Consequently, the collection of user personal data has become more aggressive and sometimes even intrusive [31, 35], raising a huge public debate around the tradeoffs between (i) innovation in advertising and marketing, and (ii) basic civil rights regarding privacy and personal data protection [57, 53]. These increasing privacy concerns, drew the attention of a significant body of research, which studied users' privacy loss in conjunction to existing user tracking techniques [1, 22, 61, 17, 54], and proposed various defense mechanisms to the users [60, 65, 66]. However, with all this pervasive ad-related user tracking and data collection, an outstanding question still remains unaddressed: *What is the monetary cost of the users' collected private data?* Or, in other terms, *how much do advertisers pay for users' privacy?*

Despite the importance of these questions, it is surprising how little is known about them There exist several reports about the *average* revenue per user (ARPU) from online advertising [32, 13, 68], but ARPU, as its name suggests, is just an average. It can be calculated by dividing the total revenue of a company by the number of its monthly active users. Computing the revenue per *individual* user is a completely different business for which very little work is available.

In particular, the FDVT [14] browser addon can estimate the value of an individual user for Facebook, by tapping on the platform's ad-planner. Also, similarly to us, a method in [63] leverages the RTB protocol and specifically its final stage, where the winning bidder (advertiser) gets notified about the auction's charge price per delivered impression. These charge prices were initially transmitted in cleartext and [63] focused solely on them. However, more and more advertising companies use encryption to reduce the risk of tampering, falsification or copying from adversary competitors. This trend

renders that method inapplicable for the current and future ad ecosystem, whose majority of companies will deploy charge price encryption. In contrast to these works, our present method takes into account all the web activity of a user (not only on Facebook [14]), and all RTB traffic, i.e., both cleartext and encrypted prices.

In this paper, we develop and evaluate a first of its kind methodology for enabling end-users to estimate in real time their actual cost for advertisers, even when the latter encrypt the prices they pay. Implemented as a browser extension, our method can tally winning bids for the ads shown to a user and display the resulting amount as she moves from site to site in real time. To summarize, in this paper, we make the following main contributions:

1. We propose the first to our knowledge holistic methodology to calculate the overall cost of a user for the RTB ad ecosystem, using both encrypted and cleartext price notifications from RTB-based auctions.

2. We study the feasibility and efficiency of our proposed method by analyzing a year-long weblog of 1600 real mobile users. Additionally, we design and perform an affordable (a few hundred dollars cost) 2-phase real world ad-campaign targeting ad-exchanges delivering cleartext and encrypted prices in order to enhance the real-users' extracted prices. We show that even with a handful of features extracted from the ad-campaign, our methodology achieves an accuracy $> 82\%$. The resulting ARPU is $\sim$55% higher than that computed based on cleartext RTB prices alone. Our findings challenge the related work's basic assumption [63] that encrypted and plain text prices are similar (we found encrypted prices to be $\sim 1.7\times$ higher). Finally, we validate our methodology by comparing our average estimated user cost with the reported per user revenue of major advertising companies.

3. We implemented a prototype of our approach, where the users, by using our Chrome browser extension, can estimate in real-time, in a privacy-preserving fashion on the client side, the overall cost of their exposed private information for the advertisers. In addition, they can also contribute anonymously their impression charge prices to a centralized platform for further research.

The rest of the paper is organized as follows: Section 2 summarises various key concepts and entities of the RTB ecosystem. Section 3 presents the main challenge and thus motivation for our work – the increasing popularity of encrypted price notifications in RTB. Section 4 provides a high-level overview of our novel methodology and our price modeling engine. Section 5 presents an analysis of a year-long weblog dataset of more than 1500 users that was used to bootstrap our modeling of encrypted prices. Section 6 presents in detail the effort to model RTB charge prices by executing probing ad-campaigns. These campaigns provide ground-truth data, which then is used to train a machine learning classifier that can estimate encrypted prices in real-time at the browser of a user. Section 7 puts all the pieces together and presents results on the overall advertising cost for displaying ads to users. Section 8 covers related work and Section 9 concludes the paper.

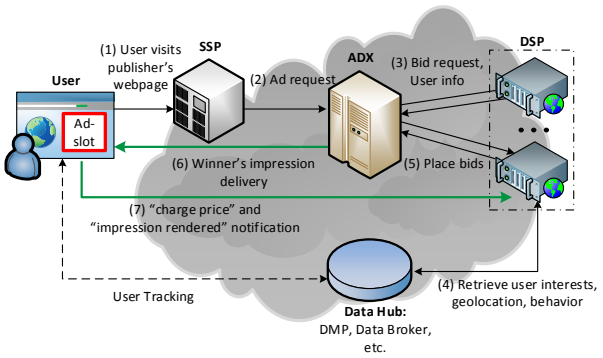

Figure 1: High level overview of the RTB ecosystem. Several entities interact with each other, exchanging user's personal data before it is finally converted to money.

2 Background on RTB

RTB accounts for 74% of programmatically purchased advertising, reaching a total revenue of $8.7 billion in the US [7]. It allows advertisers to evaluate the collected data of a given user at real-time and bid for an ad-slot in the user's display. Next, we briefly cover the most important aspects of RTB auctions, the key entities involved (§ 2.1), and how they are relevant to our study (§ 2.2).

2.1 The key players

As can be seen in Figure 1, the key roles of the RTB ecosystem include the *Advertiser*, *Publisher*, *DSP*, *Ad-exchange*, and *SSP*, which interact with each other in several ways. Note that it is very common for some (large) companies to play simultaneously different roles even inside the same auction (e.g. Google's DoubleClick Bid Manager [30] and DoubleClick for Publishers [25].

Publisher: (e.g., CNN.com) is the owner of a website, where users browse for content and receive ads (step 1). Each time a user visits the website, an auction takes place for each available ad slot. The ad impression of the winning advertiser is finally displayed in each auctioned slot of the website.

Advertiser: is the buyer of a website's ad slots. The advertiser creates ad campaigns and defines the audience that has to be targeted according to his marketing objectives, budgets, strategies, etc. In each auction, the one with the highest bid wins the ad slot and places its impression on the screen of the website's visitor.

Supply-Side Platform (SSP): is an agency platform, which enables publishers to manage their inventory of available ad slots and their pricing, allocate ad impressions from different channels (e.g. RTB or backfill in case of unsold inventory [48]) and receive revenue. SSP is also responsible for interfacing the publisher's side to multiple ad-exchanges (step 2) and aggregate/manage publisher's connections with multiple ad networks and

buyers. In addition, by using web beacons and cookie synchronization, SSPs perform user tracking in order to better configure their ad slots' pricing and achieve as many re-targeting ads as possible and thus higher bids [45]. Popular vendors selling SSP technology include: OpenX, PubMatic, Rubicon Project, Right Media.

Ad-exchange (ADX): is a digital, real-time marketplace that, similar to a stock exchange, enables advertisers and publishers to buy and sell advertising space through RTB-based auctions. ADX is responsible for hosting an RTB-based auction and distribute the ad requests along with user information (i.e. browsing history, demographics, location, cookie-related info) may own, among all the interested auction participants (step 3).

Typically these auctions follow the second higher price model (i.e. Vickrey auctions) [80], thus, the charge price for the winner of the slot is the second highest submitted. Its main advantage is that it forces all bidders to have their bids truly reflect on what they think the value of the specific ad slot should be based on the conjoined user info. After the auction, the winning impression is served to the user's display within 100 ms of the initiating call (step 6) and the winning bidder is notified about the final charge price. Popular ad-exchanges include: DoubleClick, MoPub, and OpenX.

Demand-Side Platform (DSP): is an agency platform, which employs decision engines with sophisticated audience targeting and optimization algorithms aiming to help advertisers buy the best-matched ad slots from ADXs in a simple, convenient and unified way. DSPs retrieve and process user data from several sources (step 4) such as ADXs, Data Hubs, etc. The result of this processing is translated to a decision in practice: *Does it worth to bid for an ad slot for this user, and how much?* If the visitor's profile matches the audience the advertiser has focused his ad campaign on, the DSP will submit to the ADX the impression and a bid in CPM (cost per 1000 impressions [40], typically in USD or Euro) on behalf of the advertiser (step 5). Popular DSPs include: MediaMath, Criteo, DoubleClick, AppNexus and Invite Media.

User Data Hub or Data Exchange Platform (DXP): is a centralized data warehouse such as a Data Management Platform (DMP) [47, 8] or a Data Broker [51] which aggregates, cleans, analyzes and maintains user private data such as demographics, device fingerprints, interests, online and offline contextual and behavioral information [43, 42]. These user data are typically aggregated in two formats: 1) a full, audience user profile for offline analytics and data mining, 2) a run-time user profile, optimized for real-time requests such as RTB queries from DSPs, before submitting their bids to ADXs [19, 44] (step 4).

Such user profiles are sold to ad entities [4] because they increase the value of an RTB inventory by enabling a more behavioral-targeted advertising (2.7× more effective than non-targeted advertising [5, 83]). In fact, Data Hubs are considered the core component of the digital ad-ecosystem as they perform the attribution and labeling of users' data and create groups, namely *audience segments*, which are useful (i) to the publishers for their customer understanding, (ii) to the SSPs for retrieving more re-targeted ads and (iii) to the DSPs for feeding their bid decision engines. Furthermore, quality scores are impartially assigned to users' private data based on the success of the ad campaigns they were used, thus driving the bid prices of future ad-campaigns. Notable Data Hubs include: Turn, Adobe, Krux, Bluekai, Lotame.

Winning Price Notification URLs

(A) cpp.imp.mpx.mopub.com/imp?ad_domain=amazon.es&
ads_creative_id=ID&**bid_price=0.99**&bidder_id=ID&...
bidder_name=..&**charge_price=0.95**&country=..&...
currency=USD&latency=0.116&mopub_id=ID&pub_name=..

(B) tags.mathtag.com/notify/js?exch=ruc&**price=B6A3F3C19F50C7FD**&...
3pck=http%3A%2F%2Fbeacon-eu2.rubiconproject.com%2F
beacon%2Ft%2Fce48666c-6eb4-46db-b0e9-6f4155eb557d%2F

(C) adserver-ir-p.mythings.com/ads/admainrtb.aspx?googid=ID&
width=300&height=250&...&cmpid=ID&gid=ID&mcpm=60&...
rtbwinprice=VLwbGQANXi4K21KFAAAm2ziqnOS_O5oNkFuuJw&..

Table 1: Examples of (A) cleartext and (B, C) encrypted RTB price notifications. "ID" is typically a hexadecimal number.

2.2 RTB price notification channel

When an ADX selects the winning bid of an auction, the corresponding bidder must be notified about its win to log the successful entry and the price to be paid to the ADX. One could implemented this notification in two ways: (i) with a server-to-server message between ADX and DSP, (ii) with a notification message conjoined with the price, passed through the user's browser as a callback to the DSP.

The first option is straightforward and tamper-proof; no one can modify or block these messages, allowing companies to ensure that their logs are fully synced at any time. In addition, DSPs can hide information about the transactions, the purchased ad-slots and the prices paid from the prying eyes of competitors. However, DSPs do not have any indication of the delivery of each ad, in order to inform their campaigns and budget.

Instead, the second option not only can ensure the DSP that the winning impression was indeed delivered (the callback is fired soon after the impression is rendered on the user's device), but also gives the opportunity to drop a cookie on the user's device. Therefore, the second option is the dominant one in the current market: the ADX piggybacks a notification URL (nURL) in the ad-response, which delivers to the user the winning impression and the ad (steps 6 and 7 in Figure 1). This nURL includes basically the winning DSP's domain, the charge price, the impression ID, the auction ID and other relevant logistics (see Table 1 for some examples). In this present work, we study such nURLs and the prices embedded in them, as well as how they associate with the users' browsing behavior and other personal information.

3 User Data Cost from RTB Prices

A user's personal data and preferences are at the core of contemporary digital advertising. AdTech companies spend a lot of effort to aggregate such data, build user profiles and ad-audiences, and sell them to 3rd parties.

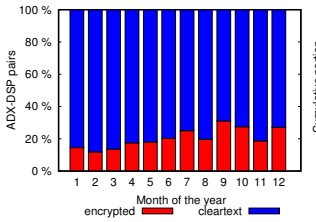

Figure 2: Percentage of encrypted and cleartext pairs of ADX-DSP through time (2015).

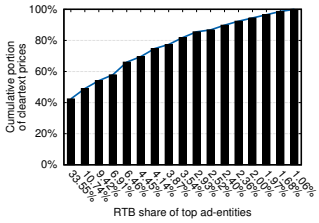

Figure 3: Cumulative portion of cleartext prices found with respect to advertising entities' portion of RTB prices found.

The main motivating question of this work is how much advertisers pay for user's personal data while delivering ads to them.

3.1 Encrypted vs. cleartext prices

Nowadays, charge prices in nURLs can be either in cleartext or encrypted form (see examples in Table 1). Cleartext prices captured at the user's browser can be easily tallied to estimate the total cleartext cost. The same does not apply for the encrypted prices though. The 28-byte encryption scheme, that is used in most of the cases, cannot be easily broken [28].

Related studies [63] assumed that encrypted prices follow the same distribution as cleartext. Indeed, one may argue that the price encryption is to avoid tampering of reported prices, and thus encrypted prices probably follow the cleartext price distribution. This could be the case for a portion of them. However, it can also be that an encrypted price may be a sign of a higher value that the DSP wants to hide, for several reasons: aggressive re-targeting because of user's previous incomplete purchases, targeting users with higher spending habits, or users with specialized needs (e.g., sensitive products, expensive drugs, etc.). Therefore, an encrypted price may be used by a DSP to reduce transparency over its bidding strategies, or possible special knowledge it may have about a specific user, thus preventing an external observer or competitor from assessing its bidding methods and ad-campaigns.

In our study, we do not adopt the above assumption and we handle encrypted prices separately. This allows us to account for any potential difference in their distribution from cleartext prices.

3.2 Encrypted prices on the rise

Encryption is a regular practice in desktop RTB ads (∼68% as reported in [62] with major supporters being DoubleClick, RubiconProject and OpenX). By analyzing a weblog of 1600 real mobile users (see Section 5), we detected a smaller portion in mobile RTB ads (∼26%). However, we found that the percentage of ADX-DSP pairs using encrypted price nURLs was steadily increasing through time (Figure 2), which means that more and more mobile advertising entities have started using nURLs with encrypted prices.

Notation	Definition
V_u	Total cost of user u
C_u, E_u	Sum of cleartext, encrypted prices of user u
SC_u, SE_u	Set of cleartext, encrypted price nURLs of user u
F_i	Vector of features for a price nURL i
$S_i \subseteq F_i$	Core features selected to represent nURL i
$ESe(S_i)$	Estimated encrypted price based on vector of features S of price nURL i

Table 2: Summary of notations.

In addition, we found that the mobile advertising entities with the larger RTB shares, deliver the highest portion of cleartext prices as well (Figure 3). For example, MoPub and Adnxs, the two leading ad-entities in our dataset, are responsible for 33.55% and 10.74% of the overall RTB ads detected, respectively (x-axis). They are also responsible for 45.40% and 5.45% for the cleartext prices detected, respectively (cumulatively in y-axis). This means that if these two companies flip their strategy from cleartext to encrypted, the overall ratio of encrypted and cleartext prices will significantly change overnight, and an external observer or the involved user will lose a significant portion (∼51%) of information.

Given these trends in mobile and desktop, we expect that in the near future RTB auctions will dominate, and many of the ad-entities will use encryption to deliver their charge prices. Our methodology anticipates these trends and promotes better transparency in online advertising and usage of user personal data. This methodology allows end-users to accurately estimate on their browser, at real-time, their average ad-related cumulative cost, even when the ad charge prices are encrypted.

4 Methodology

In this section, we describe our proposed methodology, with which a user u can estimate in real time the accumulated cost V_u for the ads she was delivered while browsing the web (§ 4.1). Table 2 presents the notations we use through this section. Following this methodology, we design our system based on two main components: (i) a remote *Price Modeling Engine* (§ 4.2) and (ii) a browser extension, namely *YourAdValue* (§ 4.3). Figure 4 presents a high level overview of our proposed methodology.

4.1 Overall cost of the user's data

The overall ad-cost of the user for time period T is the sum of charge prices the advertisers have paid after evaluating her personal data they have collected and delivering ads to her device. Specifically, this overall value is the sum of both her cleartext $C_u(T)$ and encrypted $E_u(T)$ prices and can be stated as:

$$V_u(T) = C_u(T) + E_u(T) \quad (1)$$

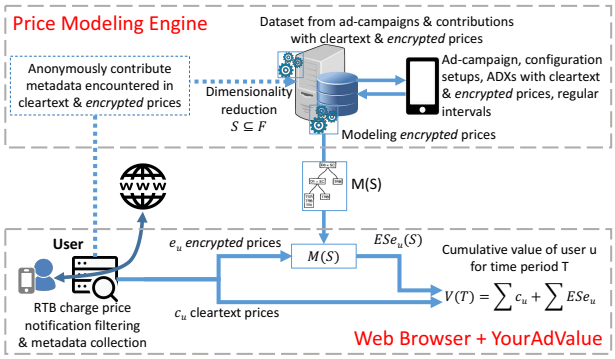
Price Modeling Engine
Dataset from ad-campaigns & contributions with cleartext & encrypted prices
Anonymously contribute metadata encountered in cleartext & encrypted prices
Dimensionality reduction
S ⊆ F
Ad-campaign, configuration setups, ADXs with cleartext & encrypted prices, regular intervals
Modeling encrypted prices
M(S)
User
e_u encrypted prices
M(S)
ESe_u(S)
Cumulative value of user u for time period T
V(T) = Σ c_u + Σ ESe_u
RTB charge price notification filtering & metadata collection
c_u cleartext prices
Web Browser + YourAdValue

Figure 4: High-level overview of our method. The user deploys a web browser extension on her device, which calculates at real-time the total cost paid for her by advertisers. In case of encrypted prices, it applies a decision tree model derived from the PME and probing ad-campaigns.

The cleartext prices of a user can be aggregated in a straightforward fashion, thus producing the ad-cost for user u over such prices:

$$C_u(T) = \sum_i c_i, \ where \ \ i \in SC_u(T) \qquad (2)$$

On the other hand, the calculation of the aggregated $E_u(T)$ of the encrypted prices for the same user is not easy. The actual price values e_i are hidden and therefore need to be estimated. To achieve that, we leverage the metadata of each charge price in the user's set SE_u(T) of encrypted price notifications. Such metadata may include: time of day, day of week, size of ad, DSP/ADX involved, location, type of device, associated IAB, type of OS, user's interests, etc. All these metadata can compose a feature vector F_i that characterize the value of a specific charge price e_i, in nURLi.

In order to estimate each encrypted notification price i, we built a machine learning model, which receives as input the feature vector F_i (or a subset $S_i \subseteq F_i$), extracted from the nURLi, and estimates a charge price $ESe(S_i)$ for the encrypted price e_i. Then, we can aggregate the estimated encrypted prices for user u:

$$E_u(T) = \sum_i ESe(S_i), \ where \ \ i \in SE_u(T) \qquad (3)$$

4.2 Price Modeling Engine

Price Modeling Engine (PME), the main component of our approach, is a centralized repository responsible for the estimation of encrypted prices. In order to achieve this, the PME requires a sample of charge price data and associated features to train a machine learning model. This component is designed to incorporate data such as offline weblogs (see Section 5), or online anonymous contributions (anonymized features and charge prices) from participating users, similarly to other works on online data and ad transparency (e.g., Floodwatch [77]). Using such data, the PME can re-train the computed model at any time. To assess the difference between cleartext and encrypted price distributions in the wild and fine-tune the training model, the PME can run small "probing ad-campaigns" to collect ground truth of real charge prices from both encrypted and cleartext formats.

Considering all possible metadata available, i.e. auctions' metadata and users' personal data, they amount to hundreds of data points per price to be fed into the modeling engine, increasing excessively the computation cost. Additionally, if all data points were to be used in the probing ad-campaigns, they would render such campaigns infeasible, money-wise. In order to run effective and efficient ad-campaigns, and enable the price model to train without high computation overhead, the PME performs dimensionality reduction of the extracted metadata (F) into a subset $S \subseteq F$ of core features that describe well the variability of cleartext prices, and consequently of encrypted prices. This dimensionality reduction makes the probing ad-campaigns feasible by reducing by many orders of magnitude the needed features of each testing setup, and effectively the number of setups to be tested (see Section 6).

Using the collected ground truth of encrypted prices from ad-campaigns, the PME trains a machine learning model M to infer encrypted prices based on their associated subset of features S. Then, each user can apply the model M (in the form of a decision tree) locally on their browser to estimate each of her encrypted charge prices based on the matching metadata S.

In case the availability of cleartext prices is limited, the reduction step to identify important features could be hindered, but not obstructed. To mitigate this, the PME can run more probing ad-campaigns to cover extra features found in users' anonymous contributions, or that are available in professional ad-campaign planners (as in FDVT [14]). Then, the most important features can be selected based on their contribution to model the encrypted prices extracted from these campaigns.

4.3 YourAdValue browser extension

YourAdValue is a browser extension responsible for monitoring the users' nURLs and calculating locally the cumulative cost of the user at real time. To achieve this, it filters nURLs from the user's web traffic and extracts (i) the RTB auction's charge prices (both encrypted and cleartext), and (ii) metadata from each specific auction (e.g. time of day, day of week, size of ad, involved DSP and ADX, etc.) along with the personal data the user leaks while using online services (e.g. location, type of device and browser, type of OS, browsing history, etc.).

As we mentioned earlier, cleartext prices can be aggregated directly, but encrypted prices must be estimated. Therefore, YourAdValue retrieves from the Price Modeling Engine a model $M(S_i)$ that (i) includes the features S_i that need to be extracted from the collected metadata, and (ii) provides a decision tree for the estimation of an encrypted price based on these features.

By following this model, the browser extension is able to estimate locally on the client side, the value $ESe(S_i)$ of the encrypted charge prices based on the features S_i of the given nURL. After estimating each encrypted price, YourAdValue presents to the user the calculated sums

Metric	D	A1	A2
Time period	12 months	13 days	8 days
Impressions	78560	632667	318964
RTB publishers	∼5.6k/month	∼0.2k	∼0.3k
IAB categories	18	16	7
Users	1594	-	-

Table 3: Summary of dataset and ad-campaigns.

Type	Feature
Geo-temporal	Time of day, Day of week
	Location of user based on IP, # of unique locations of the user, User location history
User	Interest categories of the user, Type of mobile device, # of total web beacons detected for the user, # of cookie syncs detected of the user up to now, # of publishers visited by the user, # of total bytes consumed by the user,
	Avg. number of reqs per user for the advertiser, # of HTTP reqs of the user, Avg. number of bytes per req of user, Total duration of reqs of the user, Avg. duration per req of the user
Ad	Size of ad, ADX of nURL, DSP of nURL, IAB category of the publisher, popularity of particular ad-campaign,
	# of total HTTP reqs of the advertiser, # of bytes of HTTP req, Avg. duration of the reqs for the advertiser, # of URL parameters, Number of total bytes delivered for the advertiser

Table 4: Features extracted by summarizing data from parameters embedded in each price notification detected in the dataset for users and advertisers.

$C_u(T)$ and $E_u(T)$ along with relevant statistics and the total amount $V_u(T)$ paid by advertisers (see Section 7).

5 Bootstrapping PME

In order to assess the feasibility and effectiveness of our methodology, we bootstrap the PME and train our model on real data by collecting a year long dataset containing weblogs from 1594 volunteering mobile users from the same country. Our users agreed to use a server of our control as a proxy, thus allowing us to monitor their outgoing HTTP traffic[1]. As a result, we were able to collect a large dataset D of 373M HTTP requests spanning the entire year of 2015. Note that though our dataset consists of HTTP-only traffic, in principle our approach works with HTTPS as well, using as input the contributed data of the users as can be seen in Figure 4. Table 3 presents a summary of our dataset D. Next, we present the data collection and analysis to extract features used in the price modeling and ad-campaign planning.

[1]No personally identifiable information was gathered or used and all data used were treated anonymously.

5.1 Dataset analysis

Weblog Ads Analyzer. In order to process our dataset, we implemented a weblog advertisements analyzer, capable of detecting and extracting RTB-related ad traffic.

First, the analyzer uses a traffic classification module to categorize HTTP requests based on an integrated blacklist of the popular web browser adblocker Disconnect [15].[2] Using this blacklist, the analyzer categorizes domains in 5 groups based on the content they deliver: (i) Advertising, (ii) Analytics, (iii) Social, (iv) 3rd party content, (v) Rest. It consequently applies a second-level filtering on the advertising traffic by parsing each URL for any RTB-related parameters (like nURL). The analyzer detects nURLs by applying pattern matching against a list of macros we collected after (i) manual inspection and past papers [55, 63], and (ii) studying the existing RTB APIs [70, 27, 64, 37, 58] used by the current dominant advertising companies. From these detected nURLs, it extracts the charged prices paid by the winning bidders, after filtering out any bidding prices that may co-exist in each nURL. It also extracts additional ad-related parameters such as ad impression ID, bidder's name, ad campaign ID, auction's ad-slot size, carrier, etc.

Other operations carried out by our analyzer include: (i) user localization based on reverse IP geo-coding, (ii) separation of mobile web browser and application originated traffic based on the `user-agent` field of each HTTP request, (iii) extraction of device-related attributes from the `user-agent` field (type of device, screen size, OS etc.), (iv) identification of cooperating ADXs - DSPs pairs, leveraging the nURL used by the ADX to inform the bidder (i.e. DSP) about its auction win, (v) user interest profile based on web browsing history.

Feature extraction. There are different machine learning algorithms DSPs could use for their decision engines. They can take various features as input, each affecting differently the bidding price and, consequently, the charged price of an ad-slot. To identify such important parameters, we extracted several features from the nURLs of our dataset such as user mobility patterns, temporal features, user interests, device characteristics, ad-slot sizes, cookie synchronizations [1], publisher ranking, etc. Next, we present the analysis of the most interesting features (Table 4 presents a summary). We group them into 3 categories: geo-temporal state of the auction (§ 5.2), user's characteristics (§ 5.3), and ad-related (§ 5.4).

5.2 Geo-temporal features

An important parameter that could affect a price of an ad-impression in a RTB auction is the user's current location, information which is broadly available to publishers and trackers. Therefore, in our dataset we extract user IP address and using the publicly accessible MaxMind geoIP database [56], we map each IP to its city level. In Figure 5, which presents the min, 20th, 50th,

[2]Our analyzer can also integrate more than one blacklists (e.g., Adblock Plus' Easylist, Ghostery's blacklist, etc.)

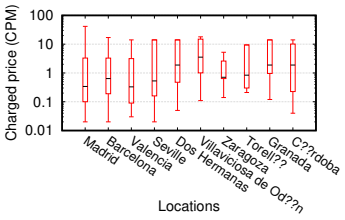

Figure 5: Distribution of charged prices for different cities (sorted by decreasing city size).

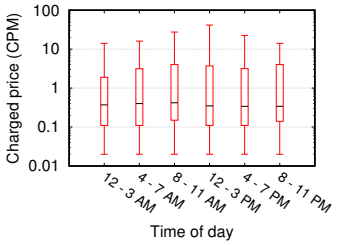

Figure 6: Distribution of charge prices for different times of day.

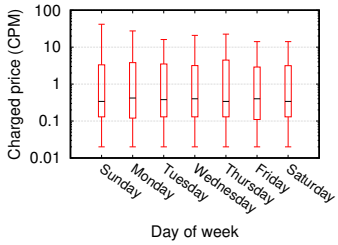

Figure 7: Distribution of charge prices for different days of week.

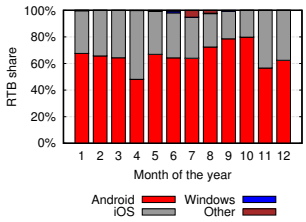

Figure 8: Percentage of RTB traffic for top mobile OSes.

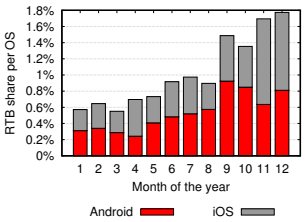

Figure 9: Percentage of RTB traffic normalized per mobile OS

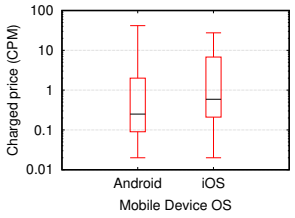

Figure 10: Distribution of charge prices for different mobile OSes.

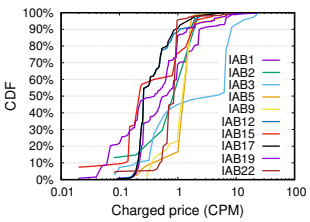

Figure 11: CDF of the generated cost by each IAB category for a 2-month period.

95th percentile and max of the charged prices, we see that although the median values are relatively lower in large cities, the fluctuation of their price values is higher.

Another important feature is time, and specifically the time of day and day of week. This is important due to the different level of attention a user may give to an ad impression and the amount of time she has to purchase an advertised product (e.g., working hours vs. afternoon's free time, or weekdays vs. weekends). From Figure 6, although the median charge prices are of similar range, the early morning hours until noon tend to have more charge prices with increased values. Similarly, from Figure 7, we see that on Sundays the maximum attracted charge prices are higher than the rest of the days.

5.3 User-related features

Device type. By parsing the `user-agent` (UA) header information, our analyzer classifies traffic and inspects the different fingerprints the UA leaks (specifications of process virtual machine (e.g., Dalvik or ART) or kernel (e.g., Darwin), operating system, browser vendor etc.) Thus, we are able to identify the type of device (PC or mobile), the different types of mobile operating systems (Android, iOS, Windows) and if the traffic was generated from a mobile application or a mobile web browser.

In Figure 8, we see the percentage of RTB traffic for the different OSes over time. As expected, Android and iOS dominate, owning the larger portions of the market through the entire year, with Android-based devices appearing in 2x times more RTB auctions. However, when normalizing this RTB share per mobile OS (Figure 9), we find that Android and iOS devices are delivered mostly equal RTB impressions, with some months Android surpassing iOS and vice-versa. Then, we extract the traffic originated from the most popular ad-entity, MoPub [59], and analyze the charged prices of the impressions rendered in the different OSes. Surprisingly, although Android-based devices are more popular, we see in Figure 10 that iOS-based devices tend to receive higher RTB prices, in median values.

Inference of the user's interest. The browsing history of a user is highly related with her interests. By monitoring the websites a user visits through time, a tracker can infer her interests, political or sexual preferences, hobbies, etc., quite accurately [6]. To enrich our set of features with the users' interests, we collect all the websites each user visits across her whole network activity. Such information is available to the RTB ecosystem as well, by using cookie synchronization [1] or web beacons [36]. To extract the interests from the visited websites, we retrieve the associated categories of content for each website according to Google AdWords [29]. Then, we aggregate across groups of categories for each user and get the final weighted group of interests for each user in the form of IAB categories [39]. Figure 11, presents for the top mobile ADX (MoPub) a distribution of the generated ad revenue for the different IAB content categories in a 2 month subset of our dataset. As expected, not all IAB categories cost the same. Indeed, there are categories that are associated with products able to draw higher bid prices in auctions like IAB-3: Business & Marketing with an average charged price of up to 5 CPM for the 50% of the cases. On the other hand, there are categories like IAB-15:Science), which are unable to draw prices higher than 0.2 CPM for the 50% of the cases.

5.4 Ad-related features

Ad-slot sizes. Some ad-entities carry in their nURLs a parameter with the size of the auctioned ad-slot. In Figure 12 we plot the popularity of each of the ad-slot sizes through time. It's interesting to see that 300x250 ad-slots (known as "MPUs" or "Medium Rectangles") dominate

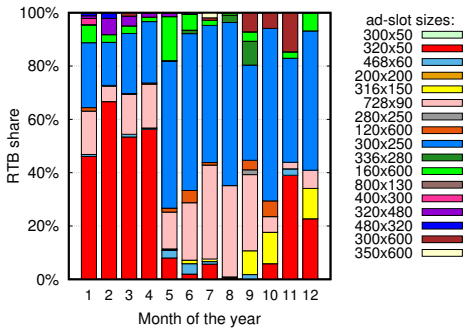

Figure 12: Ad-slot size popularity through time (sorted by area size).

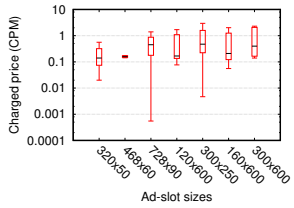

Figure 13: Distribution of the charged prices per ad-slot size (sorted by area size).

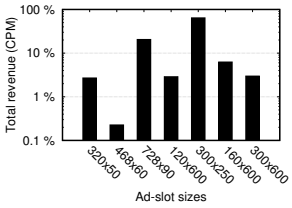

Figure 14: Accumulated revenue of the different ad-slot sizes (sorted by area size).

the dataset from May on, taking the place of 320x50 ad-slots (known as "large mobile banners"). In fact, 300x250 ad-slots have more ad content available from advertisers, so they can increase earnings when both text and image ads are enabled. In addition, we see that the 728x90 ad-slot ("Leaderboard" or "banner") is also popular. This ad-slot is usually placed at the top of a page to be seen by the user immediately when the page loads.

It is easy to anticipate, that the more space an ad-slot covers in the user's display, the higher the price will be. To verify this intuition, we isolated the traffic of an ad-entity (i.e. Turn [19]), which carries the ad-slot size in its nURLs along with the associated charged prices. Surprisingly, in Figure 13, we see that this intuition is wrong since the most expensive ad-slots for an advertiser are in fact, not the largest ones. In our dataset, we see that the two most expensive ad-slots are the MPU (300x250) and Monster MPU (300x600), with median prices of 0.47 and 0.39 CPM respectively. However, from Figure 14, the increased popularity of MPU and Leaderboard ad-slots, allows them to accumulate 64.3% and 20.6% of the total RTB revenue of Turn in our dataset, respectively. Finally, it is worth to note that our results verify past resources [16, 26] regarding the more expensive ad-slots.

5.5 Summary

In summary, by analyzing the features extracted from our offline dataset, we find that a user's location (at city level) affects the median price that advertisers would be charged, but also its variability. However, such price differences are expected to be more evident at the country-level, as shown in [63]. In addition, the days and hours that a user may not be busy (Sundays), or may offer more attention (e.g., early mornings, Mondays) lead to higher charged prices. The type of user's device also affects the charge prices but in a rather contradicting fashion: though there are more Android devices, iOS-based devices draw higher median prices. As expected, the total revenue per category of user interest (through IABs) varies a lot, with some IABs being more costly than others. Finally, the display real-estate occupied by an ad-slot does not correlate well with price. In fact, larger ad-slot sizes do not mean higher prices. As shown in the next section, these extracted features are used to plan effective ad-campaigns and model encrypted charge prices.

6 Charge Price Estimation

In order to create a model that estimates the encrypted prices detected on the user's browser and computes the total cost advertisers pay for her personal data, we need to have ground truth on charged encrypted prices. However, such dataset is not easy to acquire. One way to obtain this information is to collaborate directly with an ADX that sends such encrypted prices.[3] We assume this to be the rare case, since ADXs are unwilling to share such kind of data that may reveal bidding strategies and revenues.

To address this lack of insight on encrypted prices, we designed our system able to run small probing ad-campaigns through such ADX(s) and collect the winning prices via reports from the collaborating DSP(s). These ad-campaigns can be designed and executed with the help of a single or few DSPs any time, with little overhead and a small budget of a few hundred dollars each. In addition, they can be optimized by using a specific set of experimental setups, which cover all possible scenarios from the small parameter vector S to be kept short, efficient and cheap. Given that the prices do not change drastically over time, these campaigns can be executed every few months to collect probing data for *time-shift correction* and increased coverage of more ADXs. Besides, they can be automated and re-launched as frequently as needed, e.g., every few months or when the detected cleartext prices deviate from historical data. Having such campaigns launched from a specific location allows for more accurate and cost efficient price modeling that can be shared across all participating platform users in the same area or country.

We envision that these campaigns can be crowd-funded (e.g. like Tor Project [78], Wikipedia, WiGLE [9], etc.), thus contributing to an independent and sustainable platform that can scale better across users, countries, and ADXs covered. One may argue that these probing campaigns could pollute users' browsing with non-useful ad impressions. Therefore, they need to comply with the current standards, and if possible, consider an actual product or service. Of course, ADXs could in principle fight back and try to identify and block such campaigns, but their huge clientele combined with the low volume of such campaigns makes the detection very difficult. In the next paragraphs, we first describe

[3] We considered ADXs for encrypted prices such as DoubleClick, OpenX, RubiconProject and PulsePoint, as well as ADXs that send cleartext prices such as MoPub (the top mobile ADX).

the effort to select a subset of core features important for price modeling (§ 6.1) and how they allow us to design efficient and effective ad-campaigns (§ 6.2). Then, we provide an analysis of the data collected by two such campaigns (§ 6.3) and, finally, we describe the model that estimates encrypted prices and can be used by end-users (§ 6.4).

6.1 Dimensionality reduction of features

In order to perform ad-campaigns that are both effective and cost efficient, we need to select a subset of features $S \subseteq F$ from the available ones F that best describe the RTB prices found in weblogs such as the historic dataset D. This subset of features should explain as much of the variability of prices as possible, but should also be small. The fewer features we select as important, the smaller the cost of running probing ad-campaigns to collect representative RTB prices using these features.

To achieve this selection, we performed dimensionality reduction using all the available features (288) described in Section 5, using the cleartext prices as the target variable for optimization. First, for normalization, we applied a log transformation on the extracted cleartext prices found in D. Then, we applied a clustering of the prices into 4 classes, using an unsupervised equidistance model that finds the optimal splits between given prices using a method of leave-one-out estimate of the entropy of values in each class. Next, we filtered out features that did not vary at all (i.e., constants) or had very high variance (99%) (i.e., likely to be noise).

As a final step, dimensionality reduction (or feature selection) techniques such as PCA or random forests can be used [46]. We chose the Random Forests model (RF)[4] because it takes into account the target variable (cleartext price), it can be trained quickly on large datasets, it maintains interpretability of features and generally does not overfit the given data. In case the availability of cleartext prices is limited, the reduction step to identify important features to be used in ad-campaigns could be hindered. To mitigate this, the PME can use intermediate techniques such as high correlation filters that do not require a target variable, to eliminate features carrying similar information. We trained various RF models using subsets of semantically related features from the available feature set and the best features from each subset were selected based on their power to describe the cleartext price distribution. In summary, we grouped features in the following sets: A) time, B) http-related, C) advertisement-related, D) DSP-related, E) publisher/host interests, F) user http statistics (historical), G) user interests (historical), and H) user locations (historical). We also tried selecting representative features out of each set to create minimal combinations covering all aspects of the http-available information.

In total, we tried tens of feature subsets and combinations and evaluated them using standard machine learning metrics such as precision, recall, weighted area under the receiver operating characteristic curve (AUCROC) and out-of-bag error. Dimensionality reduction could, in principle, lead to loss of accuracy in the effort to explain price classes. However, our experimentation lead to a small subset of features with minimal loss of precision ($< 2\%$) and recall ($< 6\%$). In fact, we conclude that an optimal subset that performs very well and is small enough to allow cost efficient ad-campaigns is a set that *combines features from different groups*. In particular, (also confirmed with an ad-campaign expert) we select the following features to be used for the probing ad-campaigns described next: S={application/web-browsing, device type, user location, time of day, day of week, ad format (size), type of website, ad-exchange}.

[4] An ensemble of decision trees built using a random subset from the available features.

Filter name	Range of values (type)
Cities	city1, city2, city3, city4
Type of online interactions	Mobile in-app, Mobile web
Time of day	12am-9am, 9am-6pm, 6pm-12am
Day of week	Weekday, Weekend
Type of device	Smartphone, Tablet
Type of operating system	iOS, Android
Ad-format for smartphone	320x50, 300x250, 320x480 or 480x320
Ad-format for tablet	728x90, 300x250, 768x1024 or 1024x768
Ad-exchange	MoPub, OpenX, Rubicon, DoubleClick, PulsePoint
Categories of targeting	all IABs possible

Table 5: Basic filters used in controlled ad-campaigns in a specific country. In total, 144 experimental setups were attempted.

6.2 Ad-campaigns setup

Using the most important parameters extracted in set S, we construct various experimental setups $s \in S \subseteq F$ that can be used to deploy such ad-campaigns over a short period of time T' in selected ADXs. These setups combine different values of control variables that are important for an ad-campaign:<user location, web-interaction type, time of day, day of week, device type, OS, ad-size, ADX>. For example, an experimental setup could be this: <New York, app, 12am-9am, weekday, smartphone, iOS, 320x50, MoPub> (144 setups, Table 5). By running such controlled ad-campaigns, we can receive ground truth data about encrypted prices, thereby allowing us to train a model for such prices. Campaigns with ADXs that deliver cleartext prices also allow us to compare prices in different times and compute shifts in the price distribution due to time passed between the collection of dataset D and present time. To compensate for the loss of information from cleartext prices becoming less abundant, additional features available in professional ad-campaign planners (such as in FDVT [14]) could be used in the future to enhance the setups tested. With the results of these campaigns (in essence, charged prices for RTB ads that fulfil a given setup s), the PME can train a model to estimate the cost of new ads with a given setup s' close, or equal, to one tested, i.e. $s' \sim s \in S$.

Number of required ad-campaigns. Two important questions while designing these probing ad-campaigns are how many ad-campaigns and impressions per campaign are enough to acquire a good approximation of the mean of price distribution per campaign. For this, we analyzed the ad-campaigns found for MoPub in D. We identified 280 such campaigns in 2015, with mean and standard deviation of charge price of $m = 1.84$ and $std = 2.15$ CPM, respectively. We use the process described in [41] and the next formulation to compute d, the expected error on the mean, assuming a suggested number of setups n, and ignoring the finite population correction adjustment (thus assuming a more conservative approximation of n):

$$d = \frac{Z_{\alpha/2} \times std}{\sqrt{n}}, \quad (4)$$

where Z is the z-score of normal distribution. Using the 144 setups proposed, we can approximate to more than 95% CI (i.e. $\alpha = 0.05$) the mean price of campaigns observed in the wild, assuming a margin of error 0.35 CPM. Also, considering the distribution of prices within the largest of ad-campaigns detected for MoPub with 1.8k impressions, we can approximate to 95% CI the mean price of a campaign, assuming a margin of error 0.1 CPM and minimum of 185 impressions per campaign.

6.3 Ad-campaigns analysis

Using these as guidelines, we executed two rounds of different ad-campaigns to collect data on prices (Table 3). Note at this point, that our ad-campaigns advertised a real service in order to avoid polluting the users with meaningless impressions.

Dataset collected. The first round ($A1$) was executed for 2 weeks in May 2016 and utilized the 4 ADXs mentioned earlier that encrypt price notifications and targeted publishers of many IAB categories. The second round ($A2$) was executed with the same experimental setups as $A1$ during June 2016, but in this case the DSP was instructed to use only MoPub, while still targeting similar IAB categories of publishers. In both campaigns, the DSP was given an upper bound on the bidding CPM price, to safeguard that the allocated budget will not be consumed very quickly. On the other hand, the DSP was instructed to bid in a dynamic manner, as low or high as needed to get the minimum of impressions delivered for the various experimental setups we requested. Overall, we managed to receive for all setups, over 600k impressions displayed with encrypted price notifications to more than 200 publishers, and over 300k impressions with cleartext price notifications to more than 300 publishers, reaching audiences of 6 IAB categories common to both price notification types.

Cost paid vs. IAB category. In Figure 15, we compare the overlapping IAB categories of the RTB impressions we took from (i) the set of encrypted prices from the ad-campaign on four ADXs in $A1$, (ii) the set of cleartext prices from the ad-campaign on MoPub ($A2$), a (iii) the 2 months Mopub subset of D. Note that in some cases, the results from the dataset D vary more than in the ad-campaigns. This is to be expected, as the dataset includes prices from numerous DSP-ADX pairs for many ad-campaigns running in parallel in the duration of a year, whereas our two ad-campaigns are more targeted to specific DSP-ADX pairs.

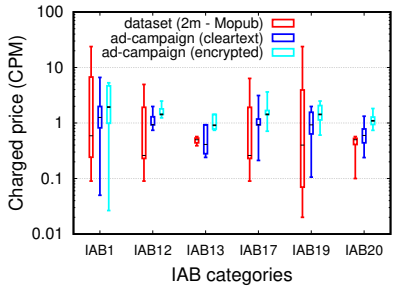

Figure 15: Comparison of CPM costs for the different IAB categories in our dataset and the 2 probing ad-campaigns.

Regarding the cleartext prices of different IAB categories, although the median prices are usually in the same order of magnitude, they are higher in the case of the recent ad-campaign contrary to the 2 month dataset. We believe that this difference is due to the time shift between the dataset collected in 2015 and the ad campaigns performed in 2016. In addition, we see that the median price is always higher in case of encrypted prices ($A1$), compared to the cleartext prices of the second ad-campaign ($A2$), and the dataset D.

6.4 Encrypted price modeling

Using the ground truth data collected from the first round of ad-campaigns with various parameters within the subset of features S, we trained a machine learning classifier to predict values of encrypted prices. As a first step, we performed similar preprocessing as earlier for the cleartext prices (normalization and clustering to 4 classes of well balanced groups). Next, we trained a RF model to predict the class value of an encrypted price, based on the available parameters S. Using features such as city of user, day of week and the time the ad was delivered, ad size, mobile OS of the user's device, IAB category of the publisher, ADX used and device type, our classifier can achieve 82.3% accuracy and 0.96 AUCROC. When the publisher used is also taken into account in the model, the performance of the classifier increases to 95%, and 0.99 AUCROC. However, this is classic overfitting and we should caution that the publishers used in the ad-campaigns are just a subset of the thousands of possible publishers that can be found in real weblogs. Therefore, we chose to use the model without the publisher as part of its input features. Next, this model was used for the estimation of the encrypted prices of nURLs found in the weblogs of each user in D, given the matching parameter values from $S \subseteq F$.

7 User Cost for Advertisers

The previous sections allowed us to: (1) bootstrap our price modeling engine from existing user weblogs, so

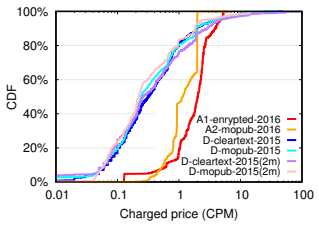

Figure 16: Comparison of price distributions between cleartext and encrypted, for different time periods and datasets (D vs. $A1$ and $A2$).

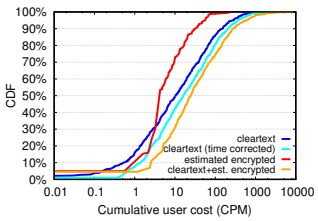

Figure 17: Cumulative CPM paid per user in our year long dataset.

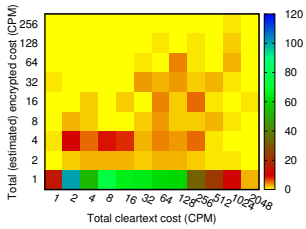

Figure 18: Total cleartext vs. total estimated encrypted cost of each user in D (color per cell indicates number of users).

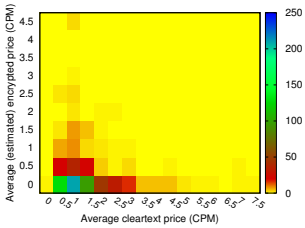

Figure 19: Average cleartext vs. average estimated encrypted price per impression of each user in D.

that we find the important features describing well the observed RTB cleartext prices, (2) using these important features, run probing ad-campaigns with ADXs that send encrypted price notifications, so that we collect ground truth on such prices from performance reports delivered to us, (3) using such ground truth, train a machine learning model to estimate the price of new RTB notifications sent in encrypted form. We are now ready to study the overall cost advertisers paid for each of the users in our dataset D, who received cleartext and/or encrypted prices in nURLs of delivered ads.

7.1 Encrypted vs. cleartext price distributions

The work in [63] assumed that encrypted prices follow the same distribution with cleartext prices. To verify the validity of such assumption, we plot the distributions of both encrypted and cleartext charge prices we got from the two ad-campaigns we performed. Interestingly, in Figure 16, we see that the price distribution of *encrypted prices* from $A1$ is distinctly different and of *higher median value* ($\sim 1.7\times$) *than cleartext prices* of $A2$.

In addition, we study the distributions between different time periods and ADXs to extract important lessons. First, we see that the cleartext price distribution of MoPub (2015) is similar to all ADXs sending cleartext prices, either when considering a 2 month period or a full year. Hence, we can study MoPub as a representative example and extrapolate lessons for the rest of the ADXs that send cleartext prices. Second, the distribution of cleartext prices from $A2$ (MoPub) are of higher median value and can be used to establish the price shift due to time difference between the time T the dataset was collected, and T' when the campaigns were executed. In reality, this price shift can be detected evenly across multiple probing ad-campaigns (e.g., once per quarter of year).

7.2 How much do advertisers pay for users' privacy?

Equipped with our presented methodology for estimating encrypted prices, we are now ready to respond to our motivating question. Specifically, we utilize our method and compute the overall cost advertisers paid for each user in the dataset D. We also apply a time-correction coefficient on the cleartext prices using the prices from the second round of ad-campaigns. This allows us to consider the increase in cleartext prices due to time difference from the weblog collection (2015) and the ad-campaigns execution (2016).

Figure 17 presents these cumulative costs in the form of CDFs of the price distributions. As expected, we observe that the cumulative cost due to encrypted prices is still not surpassing the cleartext, since the latter is still the dominant price delivery mechanism in mobile RTB. We also note that some users are more costly than others. Specifically, the median user costs $\sim$25 CPM, and up to 73% of the users cost $<$ 100 CPM through the year for the ad ecosystem in the given dataset. On the other hand, for $\sim$ 2% of users, the advertising ecosystem spent 1000-10000 CPM for the same time period. Finally, about 60% of users had an increased average cumulative cost of $\sim$ 55% on top of their cleartext cost, due to the estimated encrypted prices. These users had a median of 14.3 CPM added to their total cost, with some extreme cases of 1000-5000 CPM.

In the previous result, we compared the distributions of encrypted and cleartext prices, but while disregarding the targeted user. In order to identify if the cost paid through encrypted prices is the same with cleartext for a specific user, we compare for each user the total costs in Figure 18 and average cost per impression in Figure 19 for each type of price. We observe that a significant portion of users ($\sim$20-25%) cost similarly for ads embedded with encrypted or cleartext prices. As expected, due to the current majority of cleartext prices in the mobile ad market, a large portion of users ($\sim$75%) have higher cumulative cost from cleartext than encrypted prices. However, a small portion ($\sim$2%) costs more (2-32$\times$) in encrypted than in cleartext form, because they were delivered mostly ads with encrypted prices. When we normalize the cumulative ad cost of user per impression delivered (Figure 19), we find that for small prices of 0-3 CPM/impression, cleartext is more dominant across users. We also find a small portion ($\sim$2%) of users who cost up to 5$\times$ more CPM/impression for the delivered ads in encrypted than in cleartext form. We anticipate this portion to increase as the encrypted notification becomes the dominant delivery of RTB prices in mobile.

7.3 Summary

By studying the overall RTB advertising cost for users in our dataset, and distinguishing the encrypted from the cleartext prices, we found that the basic assumption of related work [63] that encrypted and cleartext prices are

similar, is not valid (encrypted prices are around 1.7× higher). Furthermore, advertisers, based on users' private data, paid ∼25 CPM for delivering ads to an average user, and less than ∼100 CPM for delivering ads to 3/4 of users during a year. We also identified a small portion of outlier users (∼2%) who cost 10-100× more to the ad-ecosystem than the average user, and a similar portion that costs up to 32× more in encrypted than cleartext prices, even though encrypted prices are only a quarter of the mobile RTB ecosystem.

Validation. As an effort to validate our methodology, we can extrapolate how much users cost for the ad-ecosystem and if this estimation compares with current market numbers. For this extrapolation, we make some assumptions on how our dataset represents the overall ecosystem of users and advertisers. In particular, we assume that our average mobile user, whose annual ad-cost is in the 8-102 CPM range (25th-75th perc.), has: (1) performed 2.65 hours online daily, which is ∼83% of the average daily mobile internet usage, when considering average tablet and other mobile device usage [52], (2) performed internet activity from both mobile and laptop/desktop devices, the former traffic type being ∼51% of total internet time [12], (3) received ads in a similar fashion in both HTTP and HTTPS, the former being ∼40% of the total traffic delivered to a user [73, 21], (4) received ads over RTB, which has an overhead management and intermediaries cost of ∼55% [69], and (5) received ads in a similar fashion over RTB and traditional and other online advertising, the former being ∼20% of the total online advertising [38]. Considering these factors, the overall average user ad-cost would be in the range of $0.54-6.85, which is in the order of magnitude reported by major online advertising platforms such as Twitter (owner of MoPub, ARPU: $7-8 [32]) and Facebook (ARPU: $14-17 [13]) during the period 2015-2016.

8 Related Work

There is a plethora of papers studying the user's privacy loss along with its diffusion on the web and the tracking techniques in the wild [50, 1, 22, 61, 17, 72, 54]. There are also others proposing privacy preserving countermeasures based on either (i) randomization/obfuscation-based techniques [66, 60], where the authors aim to pollute the information trackers retrieve in order to hide the users' actual user data and interests, or (ii) anti-tracking mechanisms [65, 49], where requests to trackers are avoided or blocked. All the above studies, highlight the voracity of web entities to collect more and more data about the user and her online behavior, and an arms race between the privacy-aware users and trackers.

But how do all these trackers monetize from these data? The answer is in the advertising ecosystem, where advertisers are purchasing audiences to deliver their ad-impressions. Therefore, there are several studies focusing solely on the advertising ecosystem and how a user can preserve her privacy. For example Privad [33] is a model for privacy-preserving interest-targeted advertising. It is designed to conceal a user's activities from an advertising network, by interposing an anonymizing proxy between the browser and the ad network, allowing some trusted client software to select relevant ads locally. Unfortunately, it requires broad adoption of high-performance anonymizing proxies. On the other hand, Adnostic [79] is an architecture for interest-targeted advertising that does not require tracking. Like Privad, Adnostic uses client-based functionality to perform ad selection, but eliminates anonymizing proxies at the cost of less precise ad targeting. Recently, another method was proposed [67] to obfuscate the full identity of the user from the advertisers while browsing the web. This was achieved by introducing a Web Identity Translator (WIT) in-between the user's client and the visited websites that inject cookies to her browser. WIT is responsible for hiding the user within generic cookies that (1) match the user's inferred interests from her website visits, to allow proper ads to be delivered to the user, but (2) disallows the websites and their associated trackers to collect personal data about the user. Given that advertisers are interested in adjusting their buying strategy at real time, it is unclear if such approaches can be adapted to contemporary technologies such as RTB auctions.

The economics of private data have long been an interesting topic and attracted a considerable body of research either from the user's perspective [2, 11, 75, 71], or the advertiser's perspective [63, 24, 14, 20]. Specifically, in [2] the authors discuss the value of privacy after defining two concepts (i) *Willingness To Pay*: the monetary amount users are willing to pay to protect their privacy, and (ii) *Willingness To Accept*: the compensation that users are willing to accept for their privacy loss. In two user studies [11, 75] they measure how much the users value their own offline and online personal data, and consequently how much they would sell them to advertisers, by entering bids into an auction-based system for selling data to advertisers. In [71], the authors propose "transactional" privacy, a mechanism which allows the users to decide what personal information can be released and receive compensation from selling them.

In [63], the authors perform an extensive privacy analysis of cookie matching in association with the RTB advertising. Similar to our approach, they leverage the RTB nURL to observe the charged prices and they conduct a basic study to provide some insights into these prices, by analyzing different user profiles and visiting contexts. Their results managed to confirm that when the users' browsing histories are leaked, the charged prices tend to be increased. Similarly, in [62], the authors propose a transparency enhancing tool, which shows to the users the RTB charged price every time a RTB auction is performed. Furthermore, they collect profiled and unprofiled data from a browser extension and a crawler respectively, and they compare the RTB prices, the bidding frequency and the inter-relations among ADXs and DSPs. Contrary to our work, both studies use a dataset from (i) a small number of 100 users, (ii) over desktop, (iii) covering only one month, (iv) and based on these data, they estimate the advertising total revenues using only the cleartext prices based on the arbitrary assumption that encrypted and cleartext prices follow the same distributions.

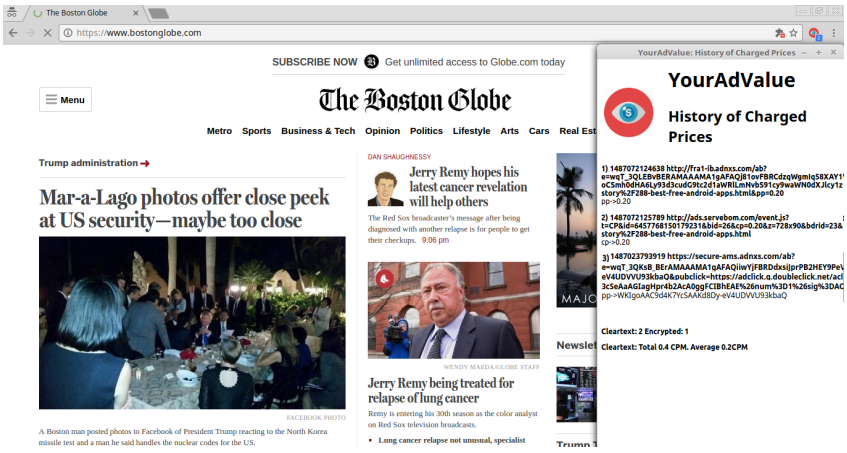
The Boston Globe
https://www.bostonglobe.com
SUBSCRIBE NOW
Get unlimited access to Globe.com today
Menu
The Boston Globe
Metro Sports Business & Tech Opinion Politics Lifestyle Arts Cars
Trump administration
Mar-a-Lago photos offer close peek at US security—maybe too close
A Boston man posted photos to Facebook of President Trump reacting to the North Korea missile test and a man he said handles the nuclear codes for the US.
DAN SHAUGHNESSY
Jerry Remy hopes his latest cancer revelation will help others
The Red Sox broadcaster's message after being diagnosed with another relapse is for people to get their checkups. 9:06 pm
Jerry Remy being treated for relapse of lung cancer
Remy is entering his 30th season as the color analyst on Red Sox television broadcasts.
Lung cancer relapse not unusual, specialist
YourAdValue: History of Charged Prices
YourAdValue
History of Charged Prices
Cleartext: 2 Encrypted: 1
Cleartext: Total 0.4 CPM. Average 0.2CPM

Figure 20: YourAdValue Chrome browser extension in use.

In [24], the authors use a dataset of several users' HTTP traces and provide rough estimates of the relative value of users by leveraging the suggested bid amounts for the visited websites, based on categories provided by the Google AdWords. FDTV [14] is a plugin that informs users in real-time about the economic value of the personal information associated to their Facebook activity. Although similar to ours, our approach works for all HTTP activity of mobile users. Furthermore, a group of journalists from Financial Times, created an interactive calculator [20] to explore how valuable the users' data are for the ad-companies. This calculator is based on the analysis of industry pricing data from a range of sources in the US.

Finally, regarding RTB auctions in particular, its rapid growth has drawn the attention of the research community, which aims to explore the economics of the RTB ad ecosystem. In [85], the authors provide an insight to pricing and an empirical analysis of the technologies involved. They use internal data of a ADX and they study its bidding behaviors and strategies. In [82], the authors propose a winning price predicting mechanism by leveraging machine learning and statistical methods to train a model using the bidding history. Their predicting approach aims to help DSPs fine-tune their bids accordingly. Though such studies help us understand some internal mechanisms of ADXs and DSPs, they are not applicable to our setting as we try to infer the cumulative ad-cost of each user based on user-related features that are measurable from the user's device over time.

9 Conclusion

In this study, we developed a first of its kind methodology to estimate how much the ad-ecosystem values the personal data of an online user. Our methodology leverages the rapidly growing real-time bidding (RTB) protocol and the new advertising model of programmatic instantaneous auctions, where the advertisers evaluate the users' collected data at real time and bid for an ad-slot in their display. Our study analyzes the RTB price notifications sent to winning advertising bidders through the users' devices. In particular, we focus on the distinction between cleartext and encrypted price notifications and how to estimate the latter. Towards this end, we train a machine learning model using as ground truth prices obtained by running our own probing ad-campaigns. We validate our methodology using a one year long trace of mobile user browsing data, as well as two real world mobile ad-campaigns and demonstrate the feasibility of computing a user's advertising value at real time.

As future work, we plan to make our prototype system (Figure 20) available for the community to test and explore its effectiveness with online users.

10 Acknowledgments

The authors would like to acknowledge the contributions and help received during the execution of this project: Prof. Vishal Misra and his team from Columbia University, for a preliminary study as part of an early DTL grant; Jose Ramon Gomez Utrilla from Telefonica's 4th Platform team, for his help with the design and execution of ad-campaigns; Costas Iordanou, for his help with the Google AdWords collection, Rafa Gross-Brown from the DTL team on the design of ads, and Xiaoyuan Yang and Martin Levi Gonzalez from the Awazza/Niji team for providing the data and help analyzing them.

References

[1] G. Acar, C. Eubank, S. Englehardt, M. Juarez, A. Narayanan, and C. Diaz. The web never forgets: Persistent tracking mechanisms in the wild. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 2014.

[2] A. Acquisti, L. K. John, and G. Loewenstein. What is privacy worth? *The Journal of Legal Studies*, 2013.

[3] C. Anderson. Free! why $0.00 is the future of business. `https://www.wired.com/2008/02/ff-free/`, 2008.

[4] BDEX - Big Data Exchange. Dmp 2.0 - introduction of the dxp. `http://www.bigdataexchange.com/dmp-2-0-introduction-of-the-dxp/`, 2015.

[5] H. Beales. The value of behavioral targeting. *Network Advertising Initiative*, 2010.

[6] P. Bernal. Our web history reveals what we think and do. shouldnt that remain private? https://theconversation.com/our-web-history-reveals-what-we-think-and-do-shouldnt-that-remain-private-50289, 2015.

[7] BI Intelligence. Programmatic advertising is under review. http://www.businessinsider.com/programmatic-advertising-under-review-2017-1, 2017.

[8] BlueKai. Data management platforms demystified. http://www.bluekai.com/files/DMP_Demystified_Whitepaper_BlueKai.pdf, 2011.

[9] bobzilla, arkasha, uhtu. Wigle: Wireless network mapping. https://wigle.net/, 2001.

[10] M. S. Brown. Your personal data's for sale: But is it accurate? http://www.forbes.com/sites/metabrown/2015/09/30/your-personal-datas-for-sale-but-is-it-accurate, 2015.

[11] J. P. Carrascal, C. Riederer, V. Erramilli, M. Cherubini, and R. de Oliveira. Your browsing behavior for a big mac: Economics of personal information online. In *Proceedings of the 22nd international conference on World Wide Web*, 2013.

[12] D. Chaffey. Mobile marketing statistics compilation. http://www.smartinsights.com/mobile-marketing/mobile-marketing-analytics/mobile-marketing-statistics/, 2016.

[13] J. Constine. Facebook crushes q2 earnings, hits 1.71b users and record share price. https://techcrunch.com/2016/07/27/facebook-earnings-q2-2016/, 2016.

[14] A. Cuevas, R. Cuevas, R. Aparicio, and J. Gonzalez. Fdvt: Data valuation tool for facebook users. In *Proceedings of the Conference on Human Factors in Computing Systems*, CHI '17, 2017.

[15] Disconnect. A faster, safer internet is one click away. https://disconnect.me/, 2011.

[16] J. Driskill. Ad size guide. http://theonlineadvertisingguide.com/ad-size-guide/300x250/, 2016.

[17] P. Eckersley. How unique is your web browser? In *Proceedings of the 10th International Conference on Privacy Enhancing Technologies*, 2010.

[18] B. Ehrenberg. How much is your personal data worth? https://www.theguardian.com/news/datablog/2014/apr/22/how-much-is-personal-data-worth, 2014.

[19] H. Elmeleegy, Y. Li, Y. Qi, P. Wilmot, M. Wu, S. Kolay, A. Dasdan, and S. Chen. Overview of turn data management platform for digital advertising. *Proc. VLDB Endow.*, 2013.

[20] E. C. Emily Steel, Callum Locke and B. Freese. How much is your personal data worth? http://www.ft.com/cms/s/2/927ca86e-d29b-11e2-88ed-00144feab7de.html, 2013.

[21] L. Encrypt. Percentage of web pages loaded by firefox using https. https://letsencrypt.org/stats/, 2017.

[22] S. Englehardt and A. Narayanan. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pages 1388–1401, New York, NY, USA, 2016. ACM.

[23] J. Fontanella-Khan. Personal data value could reach 1tn. https://www.ft.com/content/5fd7d8a8-28e5-11e2-b92c-00144feabdc0?siteedition=intl, 2012.

[24] P. Gill, V. Erramilli, A. Chaintreau, B. Krishnamurthy, K. Papagiannaki, and P. Rodriguez. Follow the money: Understanding economics of online aggregation and advertising. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, 2013.

[25] Google. Doubleclick for publishers ? small business. https://www.google.com/doubleclick/publishers/small-business/, 2016.

[26] Google AdSense. Guide to ad sizes. https://support.google.com/adsense/answer/6002621, 2016.

[27] Google Developers. Real-time bidding protocol - basics. https://developers.google.com/ad-exchange/rtb/start, 2016.

[28] Google Developers. Rtb decrypt price confirmations. https://developers.google.com/ad-exchange/rtb/response-guide/decrypt-price, 2016.

[29] Google Inc. Google AdWords. https://www.google.com/adwords/.

[30] Google Inc. Doubleclick manager. https://www.doubleclickbygoogle.com/solutions/digital-marketing/bid-manager/, 2016.

[31] A. Green. Customer data collection increased to improve customer experience, research finds. http://business-reporter.co.uk/2016/07/20/customer-data-collection-increased-improve-customer-experience-research-finds/, 2016.

[32] T. Griffin. Tren's advice for twitter. https://25iq.com/2017/01/06/trens-advice-for-twitter/, 2017.

[33] S. Guha, B. Cheng, and P. Francis. Privad: Practical privacy in online advertising. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, 2011.

[34] M. Hachman. The price of free: how apple, facebook, microsoft and google sell you to advertisers. http://www.pcworld.com/article/2986988/privacy/the-price-of-free-how-apple-facebook-microsoft-and-google-sell-you-to-advertisers.html, 2015.

[35] B. Hamilton. Google has quietly dropped ban on personally identifiable web tracking. https://tech.slashdot.org/story/16/10/22/008216/google-has-quietly-dropped-ban-on-personally-identifiable-web-tracking, 2016.

[36] W. T. Harding, A. J. Reed, and R. L. Gray. Cookies and web bugs: What they are and how they work together. 2001.

[37] IAB. Openrtb api specification version 2.4. http://www.iab.com/wp-content/uploads/2016/01/OpenRTB-API-Specification-Version-2-4-DRAFT.pdf, 2015.

[38] IHS Technology. Paving the way: how online advertising enables the digital economy of the future. https://www.iabeurope.eu/files/9614/4844/3542/IAB_IHS_Euro_Ad_Macro_FINALpdf.pdf, 2015.

[39] Interactive Advertising Bureau. IAB Tech - Lab Content Taxonomy. https://www.iab.com/guidelines/iab-quality-assurance-guidelines-qag-taxonomy/, 2015.

[40] InvestingAnswers. Cost Per Thousand (CPM). http://www.investinganswers.com/financial-dictionary/businesses-corporations/cost-thousand-cpm-2917.

[41] iSixSigma. How to determine sample size, determining sample size. https://www.isixsigma.com/tools-templates/sampling-data/how-determine-sample-size-determining-sample-size/, 2017.

[42] S. M. Julia Angwin, Terry Parris Jr. Facebook is quietly buying information from data brokers about its users' offline lives. http://www.businessinsider.com/facebook-data-brokers-2016-12, 2016.

[43] K. Kaye. Nielsen in pact to use offline data for online ad targeting. https://www.clickz.com/nielsen-in-pact-to-use-offline-data-for-online-ad-targeting/77948/, 2009.

[44] M. Kihn. Top 10 amazing secrets of dmps. http://blogs.gartner.com/martin-kihn/top-10-amazing-secrets-of-dmps/, 2016.

[45] B. Kneen. Ssp to dsp cookie syncing explained. http://www.adopsinsider.com/ad-exchanges/cookie-syncing/, 2011.

[46] Knime. Seven techniques for dimensionality reduction. https://www.knime.org/blog/seven-techniques-for-data-dimensionality-reduction, 2015.

[47] Know Online Advertising Inc. Data Management Platform DMP. http://www.knowonlineadvertising.com/programmatic-buying/data-management-platform-dmp/, 2013.

[48] Know Online Advertising Inc. Definition of backfill. http://www.knowonlineadvertising.com/advertisingdictionary/backfill/, 2013.

[49] G. Kontaxis, M. Polychronakis, A. D. Keromytis, and E. P. Markatos. Privacy-preserving social plugins. In *Proceedings of the 21st USENIX Conference on Security Symposium*, 2012.

[50] B. Krishnamurthy and C. Wills. Privacy diffusion on the web: A longitudinal perspective. In *Proceedings of the 18th International Conference on World Wide Web*, 2009.

[51] S. Kroft. The data brokers: Selling your personal information. http://www.cbsnews.com/news/the-data-brokers-selling-your-personal-information/, 2014.

[52] Leading Edge Provider. Internet trends, stats & facts in the u.s. and worldwide 2016. http://www.leadingedgeprovider.com/2016/12/internet-trends-stats-facts-in-the-u-s-and-worldwide-2016/, 2016.

[53] M. Learmonth. Online ad industry: Advertising is 'creepy'. http://adage.com/article/digital/online-ad-industry-advertising-creepy/140840/, 2009.

[54] C. Leung, J. Ren, D. Choffnes, and C. Wilson. Should you use the app for that?: Comparing the privacy implications of app- and web-based online services. In *Proceedings of the 2016 ACM on Internet Measurement Conference*, 2016.

[55] Lukasz Olejnik and Claude Castelluccia. To bid or not to bid? measuring the value of privacy in rtb. http://lukaszolejnik.com/rtb2.pdf, 2015.

[56] MaxMind Inc. Geoip databases & services: Industry leading ip intelligence. https://www.maxmind.com/en/geoip2-services-and-databases.

[57] D. Mitchell. Online ads vs. privacy. http://www.nytimes.com/2007/05/12/technology/12online.html, 2007.

[58] MoPub. Mopub openrtb 2.3 integration guide. https://dev.twitter.com/mopub-demand/marketplace-integration/openrtb.

[59] MoPub Inc. Mopub platform. http://www.mopub.com/platform/.

[60] N. Nikiforakis, W. Joosen, and B. Livshits. Privaricator: Deceiving fingerprinters with little white lies. In *Proceedings of the 24th International Conference on World Wide Web*, 2015.

[61] N. Nikiforakis, A. Kapravelos, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna. Cookieless monster: Exploring the ecosystem of web-based device fingerprinting. In *Proceedings of the 2013 IEEE Symposium on Security and Privacy*, 2013.

[62] L. Olejnik and C. Castelluccia. To bid or not to bid? measuring the value of privacy in rtb.

[63] L. Olejnik, M. Tran, and C. Castelluccia. Selling off user privacy at auction. In *21st Annual Network and Distributed System Security Symposium, NDSS, San Diego, California, USA, February 23-26*, 2014.

[64] OpenX. Rtb macros. http://docs.openx.com/Content/demandpartners/rtb_macros.html.

[65] E. P. Papadopoulos, M. Diamantaris, P. Papadopoulos, T. Petsas, S. Ioannidis, and E. P. Markatos. The long-standing privacy debate: Mobile websites vs mobile apps. In *Proceedings of the 26th International World Wide Web Conference*, WWW'17, 2017.

[66] P. Papadopoulos, A. Papadogiannakis, M. Polychronakis, A. Zarras, T. Holz, and E. P. Markatos. K-subscription: Privacy-preserving microblogging browsing through obfuscation. In *Proceedings of the 29th Annual Computer Security Applications Conference*, 2013.

[67] F. Papaodyssefs, C. Iordanou, J. Blackburn, N. Laoutaris, and K. Papagiannaki. Web identity translator: Behavioral advertising and identity privacy with wit. In *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, 2015.

[68] Patricia Gamer. Average revenue per user is an important growth driver. http://marketrealist.com/2015/02/average-revenue-per-user-is-an-important-growth-driver/, 2015.

[69] PricewaterhouseCoopers LLP,. Iab programmatic revenue report 2014 results. http://www.iab.net/media/file/PwC_IAB_Programmatic_Study.pdf, 2015.

[70] PulsePoint. Rtb implementation notes. http://docs.pulsepoint.com/display/BUYER/RTB+Implementation+Notes.

[71] C. Riederer, V. Erramilli, A. Chaintreau, B. Krishnamurthy, and P. Rodriguez. For sale : Your data: By : You. In *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, 2011.

[72] F. Roesner, T. Kohno, and D. Wetherall. Detecting and defending against third-party tracking on the web. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*, 2012.

[73] Sandvine Incorporated. Sandvine: 70% of global internet traffic will be encrypted in 2016. https://www.sandvine.com/pr/2016/2/11/sandvine-70-of-global-internet-traffic-will-be-encrypted-in-2016.html, 2016.

[74] J. Selby. The impact of big data decisions on business valuations. https://datafloq.com/read/impact-big-data-decisions-business-valuation, 2016.

[75] J. Staiano, N. Oliver, B. Lepri, R. de Oliveira, M. Caraviello, and N. Sebe. Money walks: A human-centric study on the economics of personal mobile data. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2014.

[76] Statista Inc. Premium digital advertising spending worldwide from 2015 to 2020 (in billion u.s. dollars). https://www.statista.com/statistics/237974/online-advertising-spending-worldwide/, 2016.

[77] The Office for Creative Research. Floodwatch. https://floodwatch.o-c-r.org/.

[78] The Tor Project, Inc. Tor project: Anonymity online. https://www.torproject.org/, 2002.

[79] V. Toubiana, A. Narayanan, D. Boneh, H. Nissenbaum, and S. Barocas. Adnostic: Privacy preserving targeted advertising. In *Proceedings of the 17th Annual Network and Distributed System Security Symposium*, NDSS '10, 2010.

[80] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 1961.

[81] S. Whitbeck. Rtb is growing like mad. is your mobile marketing keeping up? http://liftoff.io/rtb-growing-like-mad-mobile-marketing-keeping/, 2015.

[82] W. C.-H. Wu, M.-Y. Yeh, and M.-S. Chen. Predicting winning price in real time bidding with censored data. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015.

[83] J. Yan, N. Liu, G. Wang, W. Zhang, Y. Jiang, and Z. Chen. How much can behavioral targeting help online advertising? In *Proceedings of the 18th International Conference on World Wide Web*, 2009.

[84] P. Yared. Mobile contacts are now the real social network. https://techcrunch.com/2014/02/21/mobile-contacts-are-now-the-real-social-network/, 2014.

[85] S. Yuan, J. Wang, and X. Zhao. Real-time bidding for online advertising: Measurement and analysis. In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*, 2013.

[86] G. Zwakman, M. Aslett, J. Stamper, J. Curtis, and K. Roy. Total data market expected to reach $132bn by 2020. https://451research.com/report-short?entityId=89339&referrer=marketing, 2016.